

WHAT IS CLAIMED IS:

1 1. A routing table circuit for a router having one or more input ports for receiving
2 a message having a destination address, the router having a plurality of output ports
3 for transmitting the message, the routing table circuit comprising:

4 one or more routing table memories to store a plurality of routing table arrays,
5 said plurality of routing table arrays being arranged hierarchically in a plurality of
6 levels, wherein each of said plurality of routing table arrays is associated with a
7 predetermined subset of prefixes of said destination address, and has entries, said
8 entries including a block default route pointer field to store a block default route
9 pointer and a routing field, wherein said routing field stores a route pointer or a next
10 level pointer pointing to one of said plurality of routing table arrays in a next level;
11 and

12 a route engine to select said block default route pointer or said route pointer
13 as a return route pointer based on said destination address, wherein said return
14 route pointer determines one of said plurality of output ports for routing said
15 message, said route engine to access said plurality of routing table arrays in said
16 plurality of levels.

17 2. The routing table circuit of claim 1 wherein said plurality of routing table
18 arrays further include a tag field to designate said block default route pointer as said
19 return route pointer.

20 3. The routing table circuit of claim 1 wherein said entries include a tag field.

21 4. The routing table circuit of claim 3 wherein said block default route pointer
22 field includes said tag field.

23 5. The routing table circuit of claim 3 wherein said tag field indicates whether
24 said routing field has said route pointer.

25 6. The routing table circuit of claim 3 wherein the tag field designates whether
26 said routing field has said next level pointer.

7. The routing table circuit of claim 1 wherein at least one of said plurality of routing table arrays is associated with a table block default entry field that specifies a table block default route pointer for a subset of entries in said at least one of said plurality of routing table arrays, wherein said route engine selects said table block default route pointer as said return route pointer when a particular entry in said at least one of said plurality of routing table arrays has no particular block default route pointer or said routing field of said particular entry has no pointer.

8. The routing table circuit of claim 1 wherein at least one of said plurality of routing table arrays is associated with a table block default entry field that specifies a table block default route pointer for all entries in said at least one of said plurality of routing table arrays, wherein said route engine selects said table block default route pointer as said return route pointer when a particular entry in said at least one of said plurality of routing table arrays has no particular block default route pointer or said routing field of said particular entry has no pointer.

9. The routing table circuit of claim 1 further comprising a table default entry register to store a table block default route pointer that specifies block default route pointers for a subset of entries in an associated routing table array, wherein at least one of said plurality of routing table arrays are associated with said table block default entry register, wherein said route engine selects said table block default route pointer as said return route pointer when a particular entry in said associated routing table array has no particular block default route pointer or said routing field of said particular entry has no pointer.

10. The routing table circuit of claim 1 wherein a routing field of at least one element of said plurality of routing table arrays is populated with a pointer from the group consisting of (A) said route pointer, (B) said next level pointer, and (C) a table default route pointer.

11. The routing table circuit of claim 1 wherein a subset of said one or more routing table memories is associated with a linked list of overlapping routes, said

linked list comprising a plurality of list entries including a first list entry, wherein said first list entry includes a longest route.

12. The routing table circuit of claim 11 wherein said list entries include a prefix and a prefix length, and sorted in a descending order of said prefix length so that said first list entry includes said longest prefix.

13. The routing table circuit of claim 1 further comprising:
a route entry table that associates said return route pointer with a next hop identifier to select one of said plurality of output ports for outputting said message.

14. The routing table circuit of claim 1 wherein
said block default route pointer of a particular entry associated with said destination address is designated as said return route pointer when said route pointer associated with said particular entry is NULL,
said route pointer of said routing field is designated as said return route pointer when said route pointer is not NULL, and
said next level pointer is used to access one of said plurality of routing table arrays in said next level.

15. A routing table circuit for a router having an input port for receiving a message having a destination address, said router having a plurality of output ports for transmitting said message, said routing table circuit comprising:
one or more routing table memories to store a plurality of routing table arrays to provide a return route pointer based on plurality of said destination address, said plurality of routing table arrays being arranged hierarchically in a plurality of levels, wherein each of said plurality of routing table arrays is associated with a predetermined subset of bits of the destination address, and have a plurality of elements, said plurality of elements including a block default route pointer field to store a block default route pointer and a route pointer field to store a longest-matching route pointer and a next-table pointer field to store a next level pointer to one of said plurality of routing table arrays in a next level; and

13 a route engine to select said block default route pointer or said longest-
14 matching route pointer as a return route pointer based on said destination address,
15 wherein said return route pointer is used to determine one of said plurality of output
16 ports for routing said message, said route engine to access one of said plurality of
17 routing table arrays in said next level.

1 16. The routing table circuit of claim 15 wherein said route engine applies said
2 next level route pointer in combination with a subset of said destination address to
3 access a next element in one of said plurality of routing table arrays in said next
4 level to return said block default route pointer or said longest-matching route pointer
5 of said next element as said return route pointer.

1 17. The routing table circuit of claim 15 wherein said block default pointer of a
2 particular element is equal to NULL when said particular element has no block
3 default route pointer, and said longest-matching route pointer of said particular
4 element is equal to NULL when said particular element has no longest matching
5 route.

1 18. A router comprising:
2 a plurality of input ports that receive a message having a destination address;
3 a plurality of output ports that output said message based on a return route
4 pointer;
5 a routing table circuit to receive said message having said destination
6 address, including:
7 one or more routing table memories to store a plurality of routing table
8 arrays, wherein said plurality of routing table arrays are arranged
9 hierarchically in a plurality of levels, wherein each of said plurality of routing
10 table arrays is associated with a predetermined subset of prefixes of said
11 destination address, and have entries, said entries including a block default
12 route pointer field to store a block default route pointer and a routing field to
13 store a route pointer or a next level pointer to one of said plurality of routing
14 table arrays in a next level; and

15 a route engine to select said block default route pointer or said route
16 pointer as said return route pointer based on said destination address,
17 wherein said return route pointer determines one of said plurality of output
18 ports for routing said message, said route engine to access said plurality of
19 routing table arrays in said plurality of levels.

1 19. The router of claim 18 further comprising:
2 a route entry table that maps the selected output pointer to a next hop
3 identifier for selecting one of the plurality of output ports for outputting the message.

1 20. A method of determining an output port for a message having a destination
2 address for a router having a plurality of routing table arrays arranged hierarchically
3 including at least a first level routing table array and a second level routing table
4 array, said plurality of routing table arrays having elements, said method comprising:
5 receiving said message;
6 determining a first index into said first level routing table array based on at
7 least a first subset of said destination address;
8 accessing said first level routing table array based on said first index, said first
9 level routing table array having a first plurality of elements, where said first plurality
10 of elements include a first block default route pointer field to store a first block
11 default route pointer and a first routing field to store a first route pointer or a first next
12 level pointer, said first level routing table comprises a first tag field, said first tag field
13 to indicate whether to select said first block default route pointer or said first route
14 pointer as a return route pointer, said first tag field to indicate whether said first
15 routing field contains said first next level pointer;
16 selecting said block default route pointer as said return route pointer based
17 on said first tag field;
18 selecting said first route pointer based on said first tag field;
19 determining a second index into said second level routing table array based
20 on said first tag field and at least a second subset of said destination address;
21 accessing said second level routing table array based on said second index,
22 said second level routing table array having a second plurality of elements including
23 a second block default route pointer field to store a second block default route

pointer and a second routing field to store a second route pointer and a second next level pointer, said second level routing table array further including a second tag field to indicate whether to select said second block default route pointer or said second route pointer as said return route pointer, said second tag field to indicate whether said second routing field contains said second next level pointer;

selecting said second block default route pointer from said second level routing table array as said return route pointer based on said second tag field; and

selecting said second route pointer of said second level routing table array based on said second tag field.

21. The method of claim 20 wherein said second level routing table array is associated with a table default route field to store a table default route pointer, said method further comprising:

selecting said table default route pointer when a particular element associated with said destination address has no second block default route pointer and said second routing field contains no pointer.

22. A method of adding a route to a routing table, comprising:

receive a new route pointer associated with a new route destination address including a new route prefix and a new route prefix length;

allocate a first level routing table, and allocate a second level routing table when the new route prefix length exceeds a first predetermined threshold, the first predetermined threshold specifying a number of bits of the destination address associated with the first level routing table, the first and second level routing tables having elements, the elements including a block default route field to store a block default route, if any, and also having a routing pointer field to store a routing pointer selected from the group consisting of (A) a route pointer and (B) a next level pointer;

determine a first particular element of the first level routing table associated with the new route destination address;

store a particular next level pointer in the routing pointer field of the particular element;

determine a second particular element of the second level routing table associated with the new route destination address;

17 store the new route pointer in the routing field of the second particular
18 element; and
19 update the block default routes for a subsequent subset of elements of the
20 second routing table based on the new route prefix length.

1 23. The method of claim 22 further comprising:
2 identify among said subsequent subset of elements of the second routing
3 table a second particular element having a routing field that has a route pointer
4 having a prefix length greater than the new route prefix length;
5 obtaining a number of elements covered by said second particular element;
6 and
7 skipping memory access to said number of elements.

1 24. The method of claim 22 further comprising;
2 receive a delete-route pointer associated with a delete-route destination
3 address including a delete-route prefix and a delete-route prefix length;
4 determine a third particular element in a lowest level table of the hierarchy
5 that stores the delete-route pointer as the route pointer to provide a start element;
6 identify a block default route, if any, of the start element;
7 determine a fourth particular element in the lowest level table of the hierarchy
8 that stores the delete-route pointer as the block default route pointer as an end
9 element;
10 for a next element after the start element to the end element:
11 when the block default route for said next element equals the delete-
12 route pointer, set the block default route of said next element equal to said block
13 default route of said start element.

1 25. The method of claim 24 further comprising;
2 determining when said next element has a routing field that has a route
3 pointer having a prefix length greater than the delete-route prefix length;
4 obtaining a number of elements covered by said next element; and
5 skipping memory access to said number of elements.

26. A method of moving entries in a routing table, comprising:

- receive a first new route pointer associated with a first new route destination address including a first new route prefix and a first new route prefix length;
- allocate a first level routing table, the first new route prefix length being less than a first predetermined threshold, the first predetermined threshold specifying a number of bits of the destination address associated with the first level routing table, the first routing table having first table elements, the first table elements including a block default route field to store a block default route, if any, and also having a routing pointer field to store a routing pointer selected from the group consisting of (A) a route pointer and (B) a next level pointer;
- determine a first particular element of the first level routing table associated with the first new route destination address;
- store the first new route pointer in a routing field of the first particular element;
- update block default routes for a subsequent subset of elements of the first routing table with the first new route pointer based on the first new route prefix length;
- receive a second new route pointer associated with a second new route destination address including a portion of the first new route prefix and a second new route prefix length being longer than the first new route prefix length, wherein the second new route prefix length is greater than the first predetermined threshold;
- allocate a second level routing table, the second level routing table having second table elements, the second table elements including a block default route field to store a block default route, if any, and also having a routing pointer field to store a routing pointer selected from the group consisting of (A) a route pointer and (B) a next level pointer;
- store a next table pointer in the routing field of the first particular element;
- determine a second particular element of the second level routing table associated with the second new route destination address;
- store the second new route pointer in a routing field of the second particular element;
- store the first new route pointer in a routing field of elements of the second level routing table that precede the second particular element;

